

UNITED STATES PATENT APPLICATION
FOR
ARBITRATION BASED POWER MANAGEMENT

Inventor:

Peter R. Munguia

Docket No. 42P18957

Express Mail No: EV409356135US

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard, Seventh Floor
Los Angeles, CA 90025-1030
(408) 720-8300

ARBITRATION BASED POWER MANAGEMENT

Field Of The Invention

The present invention pertains to the field of computer systems. More 5 particularly, this invention pertains to the field of power management of devices within computer systems.

Background of the Invention

In many semiconductor devices such as embedded processors, systems-on-a-chip 10 (SOC), or other computer systems or consumer electronic devices, on-chip busses are becoming faster and wider with many associated register queues and related logic in attached unit interfaces. Split transaction capabilities on these busses have added significant depth to these queues. This is leading to a situation where on-chip busses and 15 their associated interfaces will become a significant portion of overall system power, particularly in SOC designs.

Previous efforts to reduce power consumption on busses and their associated attached units have included powering down individual units when they are not in use and reducing bus frequencies. In embedded processors and SOC designs, there may be large variances in work loads depending on the applications currently executing, time of day, 20 type of traffic, etc. The variations in workload are often under the control of various independent software drivers. These drivers typically control their particular units but do not control system resources such as busses.

In prior systems, in order to control system resources such as busses, additional communication with software that manages system resources is typically used. The

overhead and latency involved with using software to power manage system resources such as busses can be significant. This latency reduces the value of power managing system resources such as busses that need to be able to quickly change speeds based on current workload.

Brief Description of the Drawings

The invention will be understood more fully from the detailed description given below and from the accompanying drawings of embodiments of the invention which, however, should not be taken to limit the invention to the specific embodiments described, but are for explanation and understanding only.

Figure 1 is a block diagram of one embodiment of a system including several functional units coupled to a variable speed bus and an arbitration and bus clock control unit.

Figure 2 is a flow diagram of one embodiment of a method for power managing a variable speed bus.

Detailed Description

In general, one embodiment of the invention involves throttling a bus frequency based upon incoming arbitration requests from units or devices coupled to the bus.

Arbitration circuitry monitors request rates from each requestor and increases or

5 decreases the bus frequency in order to meet the bandwidth levels requested. When the bandwidth requirements increase, the bus frequency increases. When the bandwidth requirements are reduced, the bus frequency is reduced, thereby reducing power consumption.

The embodiments described herein may be implemented such that there is no

10 software intervention required to adjust the bus bandwidth once initially configured.

Software can be used to control power consumption of individual units or devices attached to the bus. In the described embodiments, the bus frequency is independently adjusted based on request rates, the requesting device characteristics, and bandwidth requirements.

15 Figure 1 is a block diagram of a system 100 including units 110, 120, and 130 coupled to a variable speed bus 155. The frequency of the variable speed bus 155 is controlled by an arbitration and bus clock control unit 150 via a clock throttling logic unit 140.

Each of the units 110, 120, and 130 are coupled to the bus 155 via bus interface logic units 112, 122, and 132, respectively. The bus interface logic units 112, 122, and 132 are further coupled to the arbitration and bus clock control unit 150. Whenever one of the units 110, 120, and 130 desires to access the bus 155, a request is made to the arbitration and bus clock control unit 150. The arbitration and bus clock control unit 150

not only performs typical arbitration functions, but also tracks how often each of the units 110, 120, and 130 requests access to the bus 155. The arbitration and bus clock control unit 150 instructs the clock throttling logic 140 to adjust the frequency of the bus 155 according to the bandwidth requirements of the units 110, 120, and 130 based on the 5 request rates of the units 110, 120, and 130. The clock throttling logic 140 derives the clock for the bus 155 from a system clock 175.

The units 110, 120, and 130 may be any of a wide variety of functional units, including, but not limited to, host processor units, video processor units, hard disk drive controller units, IEEE (Institute of Electrical and Electronic Engineers, Inc.) 1394 10 controller units, Peripheral Component Interconnect (PCI) bridge units, management processor units, and input/output controller units for slower peripheral devices.

Although system 100 is shown with only three functional units coupled to the bus 155, other embodiments are possible with a wide range of possible numbers of units coupled to the bus. Further, although system 100 is shown to be implemented on a single 15 integrated circuit die, other embodiments are possible where the bus 155 is used to coupled together discrete devices.

An example of some of the possible functions of system 100 will now be described. For this example, the bus 155 achieves a throughput of 100MB/s for every 10MHz of clock speed. The system 100 is initialized with the bus 155 at full speed, 20 which for this example is 100MHz which corresponds to 1GB/s peak bandwidth on the bus 155. Also, for this example, the unit 110 is a host processor, unit 120 is a video processor, and unit 130 is an IEEE 1394 controller.

After initialization, the arbitration and bus clock control unit 150 recognizes that there are no incoming requests from the IEEE 1394 controller, and that less than 70% of the arbitration slots are utilized. The arbitration and bus clock control unit 150 will then direct the clock throttling unit 140 to throttle the bus clock frequency down to 70MHz.

5 The recognition interval is implementation dependent and may be based upon a reasonably finite number of sequential arbitration slots. For this example, the recognition interval is 128 arbitration slots. Each arbitration slot for this example may consist of 200 bytes, which for this example would provide 500k arbitration slots per second.

For the current example, assume that at some later time the arbitration and bus
10 clock control unit 150 recognizes that the host processor unit 110 is only sustaining 10% of the arbitration slots and that the video processor 120 is using 43% of all slots. The units 110 and 120 together are using 53% of the arbitration slots with the bus running at 70MHz. The arbitration and bus clock control unit 150 then directs the clock throttling unit 140 to reduce the bus clock frequency to 40MHz, staying just ahead of the bandwidth
15 requirements.

The arbitration and bus clock control unit 150 can use average utilization over time or use other statistical methods to determine sustained bandwidth needs. The recognition interval may be short enough (perhaps 1-10us) to handle short bursts of activity such as that from a hard disk drive controller.

20 Additionally, the arbitration and bus clock control unit 150 may recognize a request from an isochronous data transfer controller, such as the IEEE 1394 controller 130. In this case, the arbitration and bus clock control unit 150 bumps up the bus clock frequency to ensure that there is adequate bandwidth to satisfy the isochronous data

transfer. The arbitration and bus clock control unit 150 will also cause the bus clock frequency to be reduced if the isochronous data transfer unit returns to an idle state.

If the system 100 is put into an idle state, where perhaps only a management processor is active and minor peripheral traffic and host processor traffic are ongoing, the bus frequency can be reduced to a minimal level (perhaps 20MHz).

The previous example described the system 100 as having the bus 155 starting out at a maximum clock frequency upon system initialization. Other embodiments are possible where the variable speed bus 155 is set at other frequencies upon system initialization. For example, in one embodiment the bus frequency is set at a frequency sufficient to only service access to a boot ROM (perhaps 2MHz). As other units become active the arbitration and bus clock control unit 150 can increase the bus frequency to accommodate the increased bandwidth needs.

Figure 2 is a flow diagram of one embodiment of a method for power managing a variable speed bus. At block 210, a number of requests are received at an arbiter from a unit coupled to a variable speed bus. Then, at step 220, a request rate is determined for the unit coupled to the variable speed bus. The clock frequency of the variable speed bus is adjusted at block 230 depending on the determined request rate. In this manner, the variable speed bus is operated at a frequency sufficient to adequately service the requesting unit, but without wasting power by operating the bus at a frequency higher than needed. The embodiment described in connection with Figure 2 may be expanded to include a number of units or devices coupled to the variable speed bus where the arbiter tracks request rates for each of the units or devices.

As seen in the discussion above, the embodiments described provide improved energy efficiency and reduce overall power requirements potentially saving other system costs associated with heat (fans, heat sinks, etc.).

In the foregoing specification the invention has been described with reference to 5 specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than in a restrictive sense.

10 Reference in the specification to "an embodiment," "one embodiment," "some embodiments," or "other embodiments" means that a particular feature, structure, or characteristic described in connection with the embodiments is included in at least some embodiments, but not necessarily all embodiments, of the invention. The various 15 appearances of "an embodiment," "one embodiment," or "some embodiments" are not necessarily all referring to the same embodiments.